EMPIRICAL CHANNEL MODELING FOR A SATELLITE AND TERRESTRIAL INTEROPERABILITY TESTBED

Eddie Hsu[1], Keith Scott, Bruce Kwan
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
Phone: +1 818 354-3212, Fax: +1 818 393-3302
{eddie.hsu, keith.scott, bruce.kwan}@jpl.nasa.gov

ABSTRACT

The emerging global information infrastructure involving a combination of satellite and terrestrial networks brings with it the increasing need in studying protocol performance within such an environment. To more accurately reflect the characteristics of a satellite channel in the particular simulation environment, an empirical model has been implemented for both the network simulator, Ns, a foundation of the Virtual InterNetwork Testbed (VINT), and the Mobile Satellite Protocol Testbed (MSPT). MSPT, a hardware-based testbed, and VINT, a software simulation system, add to the capability of a satellite and terrestrial interoperability testbed.

The empirical model uses propagation data from NASA's Ka-band Advanced Communications Technology Satellite (ACTS). The methodology for constructing the empirical model is described and is applicable to other satellite channel types. Some factors that affect link performance, such as atmospheric fading effect and station characteristics, can be adjusted in the modeling process using results from propagation research. Sample simulation traces with burst error patterns for both the MSPT and VINT/Ns are shown.

1. INTRODUCTION

The emerging global information infrastructure involving a combination of satellites and terrestrial networks brings with it the increasing need for seamless interoperation of various communications protocols within such an environment. Testbeds simulating a satellite channel using an empirical model provide a more realistic environment in which to study protocol performance than those using simpler models.

This paper describes the construction of an empirical satellite channel model for the software network simulator, Ns, a foundation of the Virtual InterNetwork Testbed (VINT) [1], and the Mobile Satellite Protocol Testbed (MSPT) [2], a hardware network emulator. VINT and MSPT form a tool set for studying protocol behavior in a global information infrastructure involving a combination of satellites and terrestrial networks. This tool set can be used as a part of a satellite and terrestrial interoperability testbed [3]. The tool set may also help the design and evaluation of protocols for communications on a planet surface and the Interplanetary Internet [4].

The MSPT is a hardware-based environment capable of accurately emulating a satellite link by injecting error bursts obtained from empirical satellite link propagation data. Its goal is to assess and compare the various transport layer protocols in a mobile satellite environment of low data rates (~2 Mbps) using actual protocol implementations. The testbed can simulate a single link at rates up to 2 Mbps while VINT provides a framework for studying the effects of scale and heterogeneity on current and future network protocols using data rates such as 155 Mbps and above. VINT allows for the composing of simulation modules for studying the interdependencies between protocols. The increasing connectivity in terms of distance and the number of users means that scale and heterogeneity become important issues in protocol mechanisms and their associated design methods [1].

The combined use of both tools facilitates the study of actual protocol implementations and complex protocol scaling and interaction in a laboratory environment. This paper describes the modeling of a Ka-band satellite link using propagation data [5] from the NASA Advanced Communications Technology Satellite [6]. The empirical model is implemented for both the MSPT and VINT.

---

[1] Corresponding author. Also with the University of Southern California.

2. THE MOBILE SATELLITE PROTOCOL TESTBED

The Mobile Satellite Protocol Testbed (Figure 1) is originally designed to evaluate the performance of transport layer protocols and Asynchronous Transfer Mode (ATM) modifications in a mobile satellite environment of low data rates (~2 Mbps). It is a hardware-based testbed capable of testing actual implementations of protocols. The goal is to determine whether standards such as TCP Extensions for High Performance (RFC-1323), TCP Selective Acknowledgment Options (RFC-2018), etc., are friendly to mobile satellite networks and to satellite networks in general. The MSPT consists of one AdTech channel emulator that adds delays and injects error to the data stream flowing through it. It can function at rates up to 2 Mbps and delays up to 8192 ms. The channel emulator can also inject Gaussian errors at specific rates. The hardware testbed, driven by processed propagation data, provides a platform that is accurate and repeatable for emulating land, aeronautical, and marine satellite channels. The testbed is also to be used in studying communications on a planet surface.

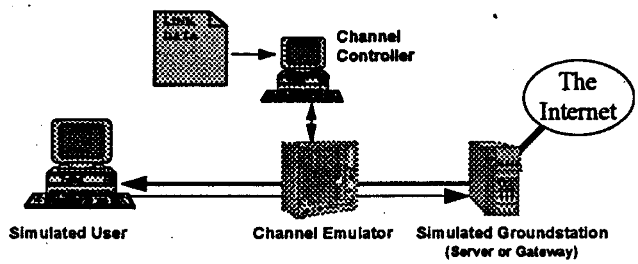

Figure 1. Configuration of the Mobile Satellite Protocol Testbed.

3. THE VIRTUAL INTERNETWORK TESTBED

The Virtual InterNetwork Testbed (VINT) is a software-based simulator system that permits the extensive study of scale and protocol interaction in current and future network protocols. VINT builds its framework upon the object-oriented network simulator, Ns [7], originally developed at the Lawrence Berkeley National Laboratory (LBNL). Nam, the LBNL network animation tool, is used for visualizing the results from Ns. VINT helps conduct protocol studies in the areas of congestion control, reliable multicast, multicast routing, dynamic topologies, and integrated services.

The system provides a framework for composing simulation modules and enables the study of complex interdependencies between multiple protocols. The system's level of abstraction, in terms of data analysis and simulations, can be varied in granularity to meet the user's needs. Scales of detailed simulations of approximately 150 nodes and 2000 links and session level simulations of approximately 2000 nodes and 8000 links have been reported [8]. Similar to the MSPT, VINT has an emulation interface that allows actual network nodes to interface with the simulator. With the available library of network topologies and traffic generators, VINT makes a useful tool in studying protocol performance over a wide range of conditions in the global information infrastructure.

4. SATELLITE LINK CHARACTERISTICS

A satellite communications link has inherent propagation delay and error characteristics that are different from those found in modern terrestrial networks.

Many satellites are located on the geosynchronous earth orbit of about 36,000 Km. The distance of a station-to-satellite-to-station hop therefore results in a propagation delay of approximately 250 ms. Satellites in lower earth orbits are placed in constellations to achieve constant coverage; therefore, the propagation delays are lower but also much more variable as routes change when a satellite traveling below the a station's communications horizon hands off a channel to another satellite still within view. Low earth orbit propagation delays range from several milliseconds to as much as 80 ms [2,9].

Signal strength from a satellite to a ground station is affected by factors such as atmospheric attenuation, ground station characteristics (antenna size, transmitter power, coding scheme), etc. [10,11]. The error pattern on a satellite link without error control coding is essentially white noise. The bit error rate for such a link is typically in the range of 10^{-1} to 10^{-5}. The use of error control coding schemes, such as Reed-Solomon, convolution, or a concatenation thereof, helps reduce the bit error rate on modern satellites links to that comparable to fiber optic cables (10^{-8} to 10^{-12}), but the use of such schemes also introduce correlated link losses. The correlated loss pattern can be captured in a multistate error model and be used to simulate satellite links in network protocol testbeds.

5. MODELING A CHANNEL USING EMPIRICAL DATA

The methodology for modeling a satellite channel using empirical data consists of the sequence of steps in Table 1. They are described in the following paragraphs. Some factors that affect link performance, such as atmospheric attenuation, transmitter power, transmitter and receiver antenna sizes, etc., can be adjusted in the modeling process. Changes in these factors can be expressed in deltas of dB to the signal levels in the fade file. The effect of the adjustments, however, will remain constant for the entire simulation session in this implementation. The adjusted signal levels are then translated to bit error rates with associated duration. The model resulting from the steps will then contain a table of bit error rates and associated duration for use in simulations.

Table 1. Steps for modeling a satellite channel.

1. Obtain the fade file for the satellite channel
2. Select a target bit error rate and an availability for the channel
3. Adjust the signal values in the fade file to meet the target availability
4. Adjust the signal values in the fade file to meet the target bit error rate
5. Adjust for other factors (e.g. atmospheric fading, antenna pointing error)
6. Translate the resulting fade file signal values to a table of bit error rates and duration

5.1 OBTAIN THE FADE FILE FOR THE SATELLITE CHANNEL

Simulating a satellite link with empirical data begins with the acquisition of a propagation data file, referred to as a *fade file*, that describes the signal levels and fade patterns on a given type of channel (e.g. Very Small Aperture Terminal with (7,1/2)-Convolution coding). A fade file from the NASA ACTS is shown in Figure 2. The signal levels in the fade file are measured in decibel (dB). The error patterns in a fade file generally are data rate independent. At higher data rates, however, other physical factors such as the longer time to re-acquire signal lock after a fade would influence the recovery period.

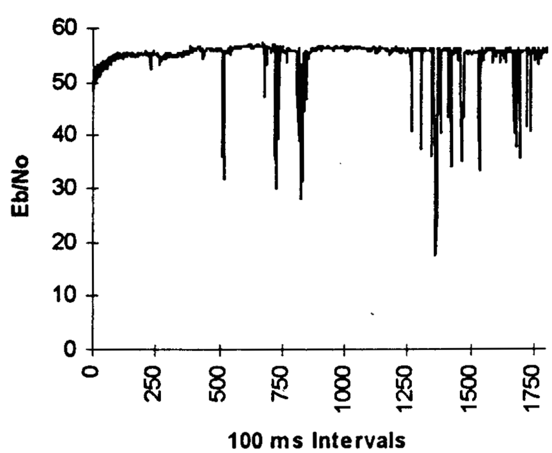

Figure 2. A propagation data file (fade file) for the NASA ACTS [5].

A goal in modeling with empirical data is to be able to accommodate a variety of ground system configurations in a particular satellite environment. While one system with large antennas and power amplifiers may be able to provide 30 dB of signal to the user, another may manage 20 dB, or even 10 dB. Therefore, while the absolute values will vary from system to system, it is the variations in the signal levels at the receiver that are of interest in the process. It is also important to extract a sufficient amount of data from a large fade file so as to reflect the statistical property of the link over a longer period of time. A way to determine an acceptable length of a fade file to use is to compare the faded or available period statistics with that of a much longer file. The fade or idle periods are categorized by their duration (e.g. multiples of 250 ms). A user may determine that a given length of a fade file is acceptable if the idle period statistics satisfactorily matches that of a much longer fade file. The goodness of selection can be determined with the help of statistical models, such as those in [5,12].

5.2 SELECT TARGET BIT ERROR RATE AND AVAILABILITY FOR THE CHANNEL

Figure 4a shows an example plot of received signal-to-noise ratio (Eb/No), in dB, from a fade file as a function of elapsed time. A 20 dB fade which drives the system below its intended operating point, *SNR**, is illustrated as the shaded portion of the figure. The shaded region is one in which the signal is too weak to be correctly interpreted by the receiver. Since the relative signal values can be used to model a variety of station configurations, the modeling steps begin with the extraction of the relative values of the signal levels. Adjustment are then applied to shift the signal levels to the concrete values that represent the particular system one intends to model.

To help perform these steps, two parameters of interest to practical systems are introduced: *target bit error rate* and *availability*. Consider, for example, a system using solely a form of Reed-Solomon encoding capable of delivering a nominal signal level (say, *SNR** of Figure 4a) to the receiver. Such a system is regarded as capable of delivering a nominal bit error rate based on the Eb/No vs. bit error rate curve for the particular configuration (Figure 3). At times when the signal falls below the nominal level, the channel can be considered faded and all data transmitted over the channel during the period can be considered lost. In addition, the fade in Figure 4a lasts 30% of the entire duration, which means the channel is available 70% of the time. Therefore we say that the system delivers a certain bit error rate, 70% of the time.

5.3 ADJUSTMENT TO MEET TARGET BIT ERROR RATE AND AVAILABILITY

To model a better system that is capable of providing a lower bit error rate for more of the time, we conceptually increase the *fade margin F* of such a system by adding the constant *F* to all the signal levels within the fade file. This in essence increases signal strength and shifts the entire curve up by *F* dB (Figure 4b). The signal profile, therefore, falls below the design threshold, *SNR**, less often, as illustrated by the smaller shaded region in Figure 4b. The stronger signal also means that the realized bit error rate is much better than the target bit error rate, since the receiver can discern the signal more clearly. A less reliable system is modeled by decreasing *F*. Other adjustment factors may then be considered before the fade file is translated into a table of bit error rates and associated duration.

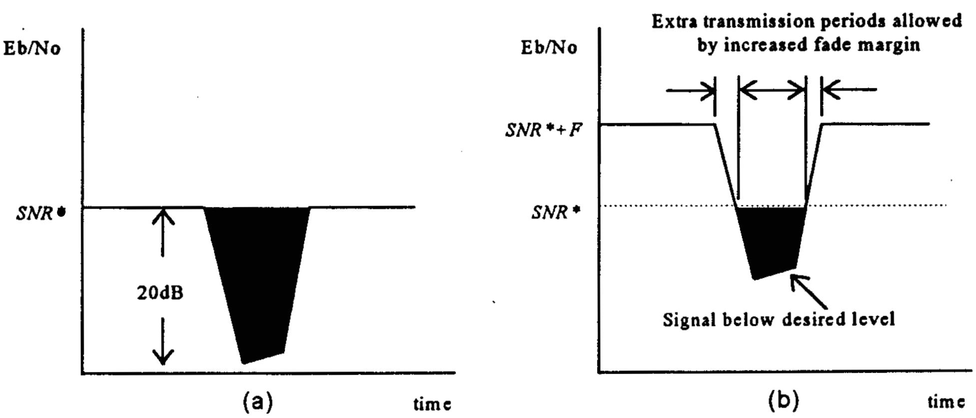

Figure 4. (a) Example SNR vs. Time Plot. A 20 dB fade is seen in the signal level at the receiver for a duration of about 30% of the time. (b) Increased Fade

Margin. A margin of *F* dB is added to the signal level. Availability is increased as less signal falls below the desired level of *SNR**.

5.4 OTHER ADJUSTMENT FACTORS

In addition to adjusting the propagation data for target bit error rate and availability, one can also take into consideration factors that can affect the reliability of a satellite link. These factors include atmospheric fading and station characteristics, and can be considered in deltas of dB to the signal levels in the fade file. [10,11] describe some of these factors and their effects on the signal level. For instance, a user may wish to reduce the signal level by 10 dB to account for rain attenuation on a Ka-band channel.

5.5 TRANSLATE THE RESULTING FADE FILE TO A PROFILE TABLE

The adjusted fade file, having taken into consideration a target bit error rate, a target availability, and other external factors, is then ready to be converted into a table of bit error rates and associated duration. The step is mapping process using the signal level vs. bit error rate curve for the particular coding scheme used (Figure 3). The resulting profile table can be used directly by the MSPT, and it needs to be converted to a multistate error model for use in the network simulator, Ns.

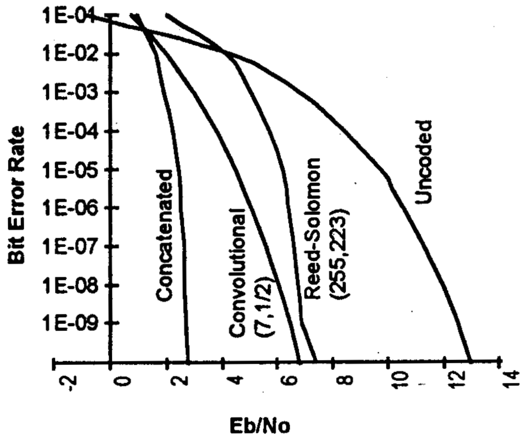

Figure 3. Signal level vs. bit error rate for various channel coding schemes, assuming Binary Phase Shift Keying (BPSK).

6. MULTISTATE ERROR MODEL FOR NS

Ns supports a wide range of models to simulate link-level errors. These error models include simpler ones based on error rates with the granularity of bytes or packets, to more complicated ones involving statistical and empirical models. For this modeling effort, a multistate error model is used. A satellite link error profile can be modeled as a transition through a series of states. The states and transitions are based on empirical data in this implementation. They may also be based on statistical link models. Each of the individual states may itself be a simple or multistate error model. It is therefore possible to construct detailed models describing various conditions that affect a link by building upon simpler forms of error models. Multistate error models, therefore, can capture more details of a satellite link.

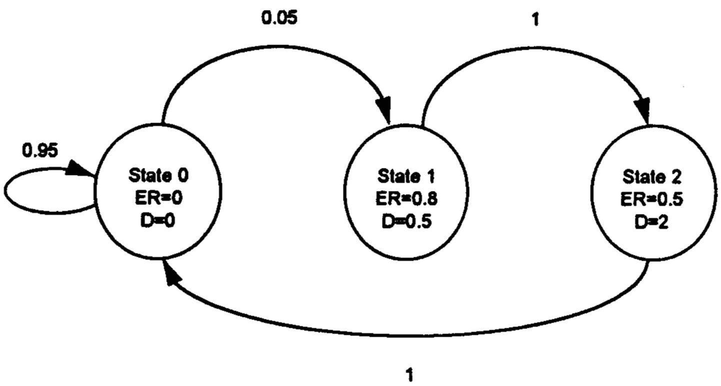

States, error rates (ER), and durations (D):

States	0	1	2
ER	0	.8	.5
D (sec)	0	0.5	2

Transition matrix and probabilities:

Current State \ Next State	0	1	2
0	.95	.05	0
1	0	0	1
2	1	0	0

Figure 6. 3-state error model transition diagram, states table and transition matrix.

The states in a multistate error model are themselves error models. Each consists of a bit error rate and, in this implementation, is assigned an associated duration based on empirical data. Departure from a state to any other state in the model depends on the probability values specified in the transition matrix (Figure 6). A state lasts for a specified duration, and the transition to the next state occurs upon the expiration of the current state's period.

In the example shown in Figure 6, there are three states (single state error models), each with an assigned error rate and duration. The duration for State 0 is 0 second; therefore, the transition out of the state is triggered immediately and is based on the probabilities in row 0 of the transition matrix. State 0 may transition to State 1 with the probability of 0.05, and it may remain in the same state with the probability of 0.95. Once a transition is made from State 0 to State 1, the model will remain in State 1 for a duration of 0.5 second. At the end of the period, the model transitions to State 2 with the probability of 1. Note that the probabilities in each row in the transition matrix add up to 1.

7. NS EMPIRICAL MODEL BURST PATTERN

To verify the burst pattern generated by the Ns empirical model, the bit error rates in the Ns model and the profile table from Section 5.5 are changed to binary numbers. A bit error rate of 0 indicates an available period, or one with signal levels at or above the nominal level; a bit error rate of 1 indicates a faded period, or one with signal levels below the nominal level. The binary scheme helps determine if the ~~the~~ empirical model properly reproduced the burst patterns.

A simple scenario is constructed in Ns where constant-bit-rate traffic is sent over a link augmented with the empirical model. For the duration of the profile table used, 174.0 seconds to the end of the last error burst, both the Ns model and the table produced 106 faded intervals. The pattern matched exactly with one shift point at 81.0 sec.. The resulting data showed that, in the Ns results, the available period lasted one extra interval before following the exact pattern again in the table.

Pearson's coefficients for the two data sets, with the Ns results shifted -5 to 5 positions relative to the profile table, are showned below (Figure 7). The figure shows the expected peaks at 1 and 0 offsets. Complete positive correlation has an r value of 0.232.

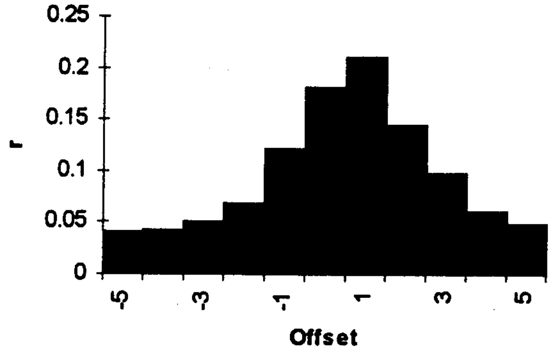

Figure 7. Pearson's coefficients comparing Ns burst patterns with that in the profile table.

8. SAMPLE SIMULATION TRACES FROM MSPT AND VINT

This section shows sample traces with burst error patterns for both the MSPT and VINT testbeds (Figure 8). The empirical model is derived from the fade file shown in Figure 2. The TCP Reno variant used in the following traces for MSPT is Linux 2.1.45 implementation with RFC-1323; the Reno variant used in Ns is a *one-way agent* that is designed to capture the essence of TCP congestion and error control behavior. Two-way agent is available in Ns and is continually being enhanced. The simulations are run for 285 seconds using 1500-byte packets at 2 Mbps. Link delays in the throughput graph (Figure 9) are in 50 ms intervals from 0 to 300 ms.

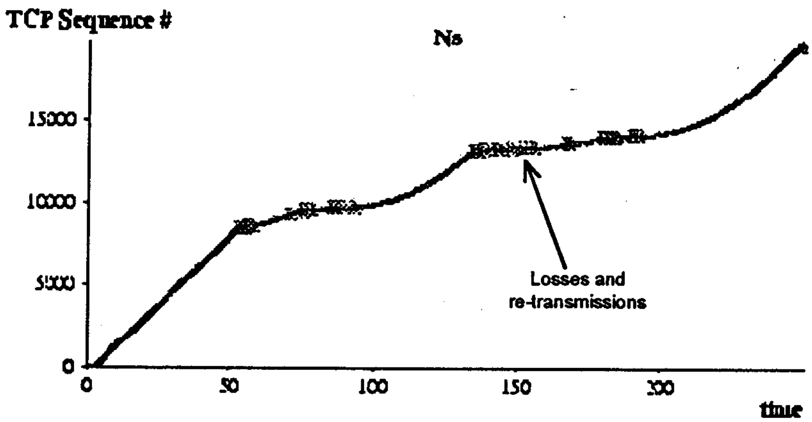

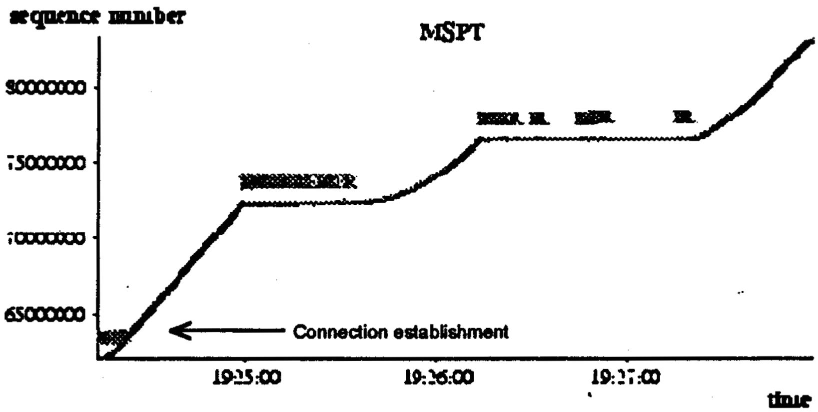

Figure 8. Sample traces showing burst error patterns. The empirical link model is derived from the propagation data shown in Figure 2. The MSPT trace also shows TCP connection establishment activity.

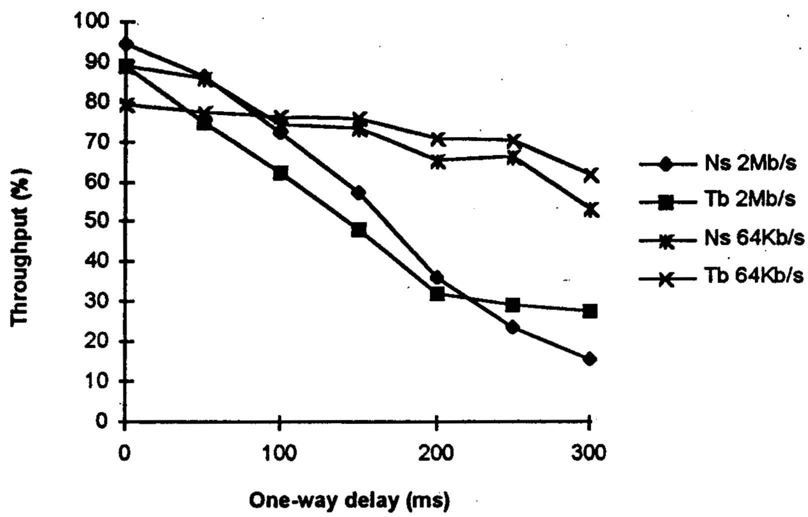

Figure 9. Throughput vs. one-way delay. Example simulations using the empirical satellite link model.

9. CONCLUSIONS

The MSPT is a low data rate hardware testbed for emulating a satellite link by injecting error bursts obtained from empirical satellite propagation data. It replicates a satellite link with great accuracy at rates up to 2 Mbps and is designed to evaluate transport layer protocols over such a link. VINT, with a foundation on the network simulator Ns, provides a framework for studying the effects of scale and heterogeneity on current and future protocols. Data rates such as 155 Mbps and above are achievable. The two tools help the study of protocol design and performance in a global information infrastructure involving satellite links.

To more accurately reflect the characteristics of a satellite channel in the particular simulation environment, an empirical model has been implemented for both the network simulator, Ns, a foundation of VINT, and the MSPT. The empirical model uses propagation data from NASA's Ka-band Advanced Communications Technology Satellite (ACTS). The methodology for constructing a satellite link error model using empirical data is described and is applicable to other satellite channel types. Statistical models can assist in determining the goodness of data selection and may be incorporated in the future into the satellite link error model.

A multistate error model is used in Ns to simulate the error characteristics of a satellite link. Sample simulation traces showing burst error patterns for both MSPT and VINT are shown. The MSPT and VINT form a tool set for studying protocol behavior in a global information infrastructure involving a combination of satellites and terrestrial networks, with future applications in areas such as communications on a planet surface and the Interplanetary Internet activity.

10. ACKNOWLEDGMENT

The activity described in this publication was carried out in part by the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

11. REFERENCES

1. Virtual InterNetwork Testbed Project. http://newweb.usc.edu/vint.
2. Keith Scott and Polly Estabrook. *Mobile Satellite Protocol Testbed.* Proceedings of the Second International Workshop on Satellite Communications in the Global Information Infrastructure. June 19, 1997.
3. Neil R. Helm. *Testbed for Satellite and Terrestrial Interoperability.* Proceedings of the Second International Workshop on Satellite Communications in the Global Information Infrastructure. June 19, 1997.
4. Adrian Hooke. *The Interplanetary Internet.* Workshop on Internets in Emerging Mixed Environments. August 5, 1998.
5. Brian Humpherys and Michael Rice. *Statistical Models for the ACTS K-Band Aeronautical Mobile Satellite Communications Channel.* Proceedings of the IEEE Vehicular Technology Conference. vol. 1, pp 46-50. May 1997.
6. D. Hoder and M. Bergamo. *Gigabit Satellite Network for NASA's Advanced Communication Technology Satellite (ACTS).* The International Journal of Satellite Communications, v.14,n.3. May/June 1996.
7. UCB/LBNL/VINT Network Simulator - Ns (Version 2). http://www-mash.cs.berkeley.edu.
8. D. Estrin. *VINT Overview.* Presentation to ISMA, UC San Diego/SDSC. May 1997.
9. Mark Allman, Dan Glover. *Enhancing TCP Over Satellite Channels using Standard Mechanisms, August 1998.* Internet Draft, draft-ietf-tcpsat-stand-mech-05.txt.
10. F. Davarian. *Ka-band Propagation Research Using ACTS.* The International Journal of Satellite Communications, v.14,n.3. May/June 1996.
11. Y. Rahmat-Samii and G. K. Noreen. *Spacecraft Antennas.* Deep Space Telecommunications Systems Engineering, Edited by Joseph H. Yuen, pp.413-447. Plenum Press. ISBN 0-306-41489-9.
12. C. Loo. *A statistical model for a land mobile satellite link.* IEEE Trans. on Vehicular Technology. vol VT-34, pp. 122-127. August 1985.